



Harmony
Email & Collaboration



YOU DESERVE THE BEST SECURITY

Protecting Against Compromised Accounts in Email

Introduction

One of the key capabilities of an email security solution is the ability to detect anomalies. When something happens that's out of the ordinary, it can be a sign that malicious behavior is afoot.

This can be referred to as account take over, or a compromised account. In an account takeover (ATO) attack, an attacker gains unauthorized access to the credentials for a user's online account. This access can then be used for identity theft, fraud, and to enable other cyberattacks, such as access to a user's corporate credentials to login and plant ransomware within the corporate network.

It's more than just taking over an account. There are real-life concerns.

In a one study, [Javelin Research](#) found that account takeover increased by a whopping 90% in 2021. These losses totaled \$11.4 billion, making up nearly a quarter of all identity fraud losses in 2021.

In a separate study, [Feedzai](#), a financial risk management company, found that account takeover was the top fraud scam, up from fourth place in 2021 and ahead of social engineering. According to Aite Group, account takeover attacks, across all industries, cost more than [\\$16 billion in losses — a 300% jump from 2020](#).

According to the 2020 Global Identity and Fraud Report by Experian, 57% of enterprises report higher fraud losses due to account takeover.

Further, according to [UK Finance](#), account takeover fraud accounts for nearly a quarter of all fraud losses.

This correlates with the data that HEC researchers see. In March alone, we saw 1,345 unique compromised accounts. **Of those, 783 began sending out phishing or spam messages. That's a 179% increase from the previous month.**

So what to do? In this whitepaper, we'll discuss how to prevent account takeover from taking control of your business.

Monitoring Account Takeover

Although phishing messages are the most common way for hackers to gain access to an account, they are far from the only method. Large, third-party data leaks like Yahoo and LinkedIn have created a market for hackers to exchange stolen passwords. Even Post-It Notes are not safe from online distribution. A breach might include passwords for one service that employees have re-used on corporate accounts. Even a breach that doesn't include raw credentials might include the personal information (street address, high school, mother's maiden name) that make it possible for attackers to gain temporary access by requesting a password change.

The Equifax breach probably contains more personal information than the average person even knows about themselves. Although anti-phishing security is important, it is only one part of the equation when it comes to defending against Account Takeover.

But it is a huge part of it. When someone takes over an account via email, they usually do one of the following actions:

- Update bank account info email to payroll
- Hijack existing legitimate conversation related to an invoice and send the attacker's bank info instead
- Create an inbox rule to hide future emails in a conversation so that the real user is unaware that anything happened
- Create a new payment request from a trusted account to a fake 3rd party service.

The idea of account take-over protection is not new. The first generation of these technologies relied on two main data monitoring techniques - monitoring suspicious logins and monitoring end-users' activities. Customers added these layers to their SSO or their SIEM and started to track geographically suspicious activities, simultaneous logins, massive downloads etc. These methods were good, but deployment was limited for several reasons:

- They require an SSO or a SIEM. Deployment of these can be a large undertaking and detecting account takeovers is only a small part of their intended use.
- SSOs only look at login events and are blind to what happens next. This isn't helpful if the suspicious activity is not the login itself, for example when the hacker was smart enough to login from the same GEO and at the right time-of-day. What if the suspicious activity is a massive download of files or sending 1,000 emails in 1 hour? SSO could not see those events and could not protect from them.

- SIEMs don't apply automated action. Some of them claim to take action, but since they are intended to be monitoring platforms, SIEMs do not truly prevent anything from happening. They don't detect in real-time so they can't auto-remediate.
- False positives are probably the most important issue, especially for SIEMs. What could be considered an anomaly may also happen regularly within the legitimate end-user activity. There are some customers that are getting 100x false-positives and therefore can't in good conscience take automated actions. After a while, end-users will stop paying attention to these alerts.

If a hacker is regularly logging into your account, wouldn't their location raise a flag? It is reasonable to assume that to detect a compromised account, you just need to keep an eye out for suspicious locations in your account history. Unfortunately, publicly available VPNs are an easy way to avoid this obvious giveaway. A competent hacker based in North Korea can appear to be from an IP address in your own town, looking as benign as a login from your local CoffeeCafe. If they've already compromised another victim, they could even stage their attack from a partner's network.

Take the story of a large hotel network based in the Northeast. As you can imagine, the hotel works with a lot of vendors. In this case, an end-user's account was taken over. The end-user worked as the Assistant Director of Finance for one of the chain's hotels. They reach out to vendor asking about an invoice:

i [redacted] ie,
 ease see the attached updated invoice for \$51,748.35 for the initial deposit, we already have \$21,117.92 on file for [redacted]. I have also attached a copy of our [redacted] help set us u
 vendor in your system, included in this email is [redacted], **Assistant Director of Finance** in case you need an more information.

There's a bunch of back and forth, until the "assistant" director of finance comes in:

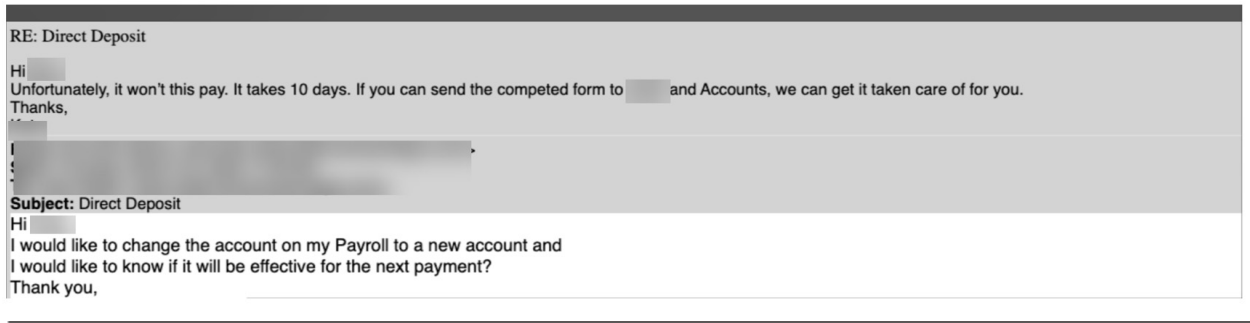
We are not in receipt of funds sent to our [redacted] bank account ending [redacted] please have your AP team to check the status of EFT transfer and advise further.

Luckily, the account was suspended before the invoice could be paid. Otherwise, that would've been \$50,000 out the door.


I went ahead and delegated [redacted] account to myself for review. There was an inbox rule set up to hide emails all emails sent/received from [redacted] st
 moved all of the hidden emails back to [redacted] nbox, removed the rule, and removed delegation access from [redacted]
 can you please review this morning and advise if anything was sent out in regards to this?

RE: SECURITY - [redacted]
 Team,
 Just spoke with [redacted] on the below. The team @ [redacted] reached out to their client through email and voicemail this morning to ensure a wire transfer was not sent o
 regarding this.
 The perpetrator(s) high jacked an existing email thread the [redacted] was on with a client and asked them to send payment out to a different account.

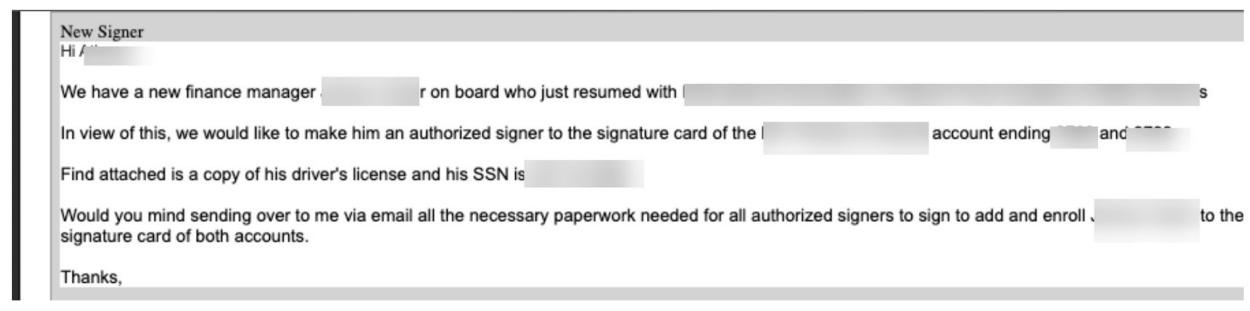
Another classic example is the Direct Deposit change. Once the hacker takes over the account, they often reach out to Payroll or HR, saying their banking information has changed. Of course, this means updating the direct deposit to the hackers' account.



In another clever example, a hacker took over an account, and then created a Twilio account. Twilio SendGrid is a service that marketing teams use to send emails. It's a legitimate service that has, in the past, [been used to send out phishing campaigns](#). To cover their tracks, the hacker created an inbox rule to delete all messages related to Twilio.



Another example showcases an attacker taking over an account and updating the authorized signer on a bank signature card.



Finally, here's a campaign in which the end-user almost paid out \$76,000.



The attacker hijacked a thread asking about a past due invoice. They inserted new banking information and attached an invoice for \$76,000;

Subtotal:	\$76,525.00 USD
Payments Applied:	\$0.00
TOTAL:	\$76,525.00 USD



Account Takeover Solutions

Now, what do you do about this? You need a solution that looks at all possible indicators of account takeovers. An organization can monitor for warning signs that an employee's account has been compromised. Some key indicators include:

- **Failed Logins:** Account takeover attacks that attempt to guess or stuff credentials on online portals can generate a large number of failed detections. Monitoring for these failed login attempts can help with detecting some types of account takeover threats.
- **User Analytics:** Users typically have certain patterns of behavior, logging in at certain times from specific places, etc. Access attempts that break these patterns of behavior can be warning signs of a compromised account.
- **Insecure Configurations:** Cybercriminals will commonly disable security controls and set up unusual configurations such as mail filtering and forwarding. These types of changes may indicate that a user account has been compromised.

You might feel that you have seen these detection methods before (and you probably have). What is unique about the Harmony Email & Collaboration solution is that those indicators, along with many more, are fed into a Machine Learning algorithm that is trained to find attacks and to filter out false positives. For example - a user that starts using a new device is something you expect to see often, and at larger organizations this event can easily happen multiple times a day. However, seeing a new device at an unexpected time of day from an unexpected geo-location and demonstrating irregular behaviors for this particular user increases the likelihood that this is indeed a compromised account. HEC also collects numerous real-world incidents of account takeover events to train our AI algorithm to find real compromised accounts at a high-accuracy while minimizing the likelihood of false-positives.

In addition, as soon as you connect the HEC platform to your SaaS, it starts scanning all historical activities. This allows us to go back to history and find the accounts that are already compromised. In a few of our customers, we found hackers that broke into the accounts of key people, such as the CFO, and waited quietly in ambush for six months waiting for an opportunity to monetize based on one of their schemes.

For each SaaS app our customers have installed, we monitor for:

- "Superman" logins (Same user, across SaaS from remote GEOs in a short time)
- First seen in a new country
- First seen with this device
- Unusual number of used devices
- Suspicious Configurations like Forward-all-to-outside-email (Office 365 + Gmail)
- Suspicious emails like an unusual number of BCC in an email (compared to other emails in the user's history)
- Login outside a defined geography
- Multiple failed login
- Same machine with login by multiple people
- Same machine sending an email and trying to login (Some hackers use hacking tools on their PC to send out phishing emails and then login to O365 from the same IP)
- Session characteristics such as length, time-of-day, behavior (For example, sending emails without reading emails, to spread an attack)
- Existence of suspicious configuration changes (For example, a Gmail rule that deletes all incoming email to hide tracks of a hacker)
- Disabling MFA
- Unusual number of 'reset password' emails received by a user
- Email rules that hide incoming emails
- Email rules that leak data to outsiders
- Phishing emails originating from an internal user
- Suspicious activity that follows the reading of a suspicious email
- Tokens with access given to low-reputation applications
- And others... (Total of over 100 indicators)

It's great to look at these indicators, and it's even better to be alerted when something happens. But taking action to ensure that a compromised internal account controlled by an attacker doesn't cause damage is key.

One way to do this is to get an alert and manually disable the user directly from the HEC dashboard. We offer that, and some companies choose to do it this way.

However, what many organizations prefer is to automatically and immediately handle such cases.

We have a workflow that automatically disables users detected as compromised and terminates all active sessions. Admins are, of course, still alerted of this and then are able to reset the password and unblock the user manually—all from the dashboard.

This is aided by our new AI engine that looks for compromised users. Instead of looking at single login parameters on their own (such as country, IP address, etc), the AI engine looks at and inspects all the parameters of login events to determine those that are done by malicious actors. This list is dynamic and constantly growing, but involves things like the IP address, browser, browser version, device, VPN brand and much more.

Login events detected by this new engine will flag their corresponding users as compromised (Critical Anomalies).

Some of these critical anomalies include:

New delete-all-emails rule

This anomaly inspects new rules configured to delete all the incoming emails. It detects potential malicious configuration to delete all the incoming emails. This behavior may indicate an account takeover.

Users Sending Malicious Emails

This anomaly is triggered when an internal user sends a phishing or spam email to internal and/or external recipients.

Move all emails to a subfolder

This anomaly inspects new rules configured to move all the incoming emails to a subfolder. It detects possible malicious configurations to move all the incoming emails to a specific subfolder. This behavior could indicate an account takeover.

Login from Malicious IP Address

This anomaly detects the compromised accounts based on the IP address from which attackers logged into Microsoft 365.

Users logging into Microsoft 365 from IP addresses detected as sources of phishing emails or from the IP address known to Check Point as malicious will be flagged as compromised.

Conclusion

Account takeover attacks can cause incredible damage. Not only can they come from anywhere, including email, but they also can lead to massive financial and reputational loss. Whether it's a man-in-the-middle attack when connecting through a shady WiFi, or an end-user that used the same credentials for another account that was then breached, these are incredibly common. Some believe Microsoft and Google already provide some detection of suspicious logins but they generally focus on the web-login itself and not on IMAP/MAPI logins, and hackers have reversed engineer that security in their lab environment and found ways to avoid its detection. Legacy email security solutions don't even claim to have an answer, they focus solely on the email flow. In the world of SaaS, a username and password is all the hackers need to take over all of your accounts. There is no need to put malware on your endpoint or get in through a Firewall. In fact, there is very little practical security after they have your username and password.

Implementing account takeover in email is key. Another key? Real-time abilities. Finding out that an account was taken over immediately after the malicious login is critical to prevent the hackers from gaining access to what they shouldn't see or using the account as jump-board to launch additional attacks.

So, now that you have accurate detection in real time, what next?

Automation is important because you often don't have time to wait for someone to review every event. It is vital to revoke the hacker's access to the account immediately before any damage is done. In addition to a low false-positive rate, your automated protection should include a method for end-users to manually remediate themselves because if it's not 'self-service' then you are relying on helpdesk calls and have users getting locked out of their accounts. That's where our automatic and instantaneous blocking comes into play.

As organizations move their email platform to SaaS based email like Office 365 and Gmail, account take-over becomes easier for hackers because all it takes is the username and password. It also becomes more valuable because all corporate data resides in the cloud, both in the emails and the other services of the suite such as OneDrive, SharePoint and Google-Drive.

A best of breed anti-phishing solution is a critical component, but is not enough because legacy email security solutions have limited success blocking current phishing attacks and because users can lose their credentials in other ways. Therefore, adding a post-attack layer to detect account takeover is a critical part of email security.

Traditional SEG-based email security platforms like Proofpoint and Mimecast cannot add this layer because they only analyze the email traffic and not the end-users' activity in the account. This is one of several reasons legacy email platforms have struggled to properly secure Office 365 and G-Suite.

An effective account takeover security layer needs to detect and respond in real-time. It must be automated and therefore it also has to be accurate and allow self-remediation for the end-users. In addition to this, it must also be able to review historical data, which can provide useful indicators of compromised accounts and give insight into what range of behavior should be considered anomalous for each end user.

HEC has a unique combination of account takeover abilities. From real-time prevention, historical breach detection, and adaptive false positive filtering, HEC provides the most advanced protection of account takeover.



Harmony
Email & Collaboration